

[0003] FIG. 1 illustrates a layout of a tape format in conformance with the Linear Tape Open (LTO) Ultrium format, which is a serpentine tape format technology. The length of the tape is divided into logical points (LPs), which define bounds of regions of the tape. The regions of LP0 to LP1 and LP6 and LP7 are unused as they are at the beginning and end of the tape, the region of LP1 to LP2 is a servo acquisition area, LP2 to LP3 is a calibration area that includes different information in the different

[0005] To determine the longitudinal position, the tape drive uses the Tape Directory which is stored either on tape or in some auxiliary cartridge memory (e.g. LTO's CM). The tape directory will typically include information at different

SECRET

5 interpolation to estimate the target's position (e.g. if we know there are 1 million
1Kbyte records in a wrap section, then interpolation would suggest that the 400,000th
record is approximately 40% of the way between the beginning and end of that wrap
section).

10 to access data. In fact, the time to wind the tape to the correct longitudinal position
can take a minute or more. Thus, the tape access time is primarily determined by the
time needed to longitudinally wind the tape to the correct position.

15 positions within the user data area, e.g., the area between LP3 to LP4. The longer the user data area or tape length, the more time required to wind or position the tape medium under the tape head.

[0008] Thus, there is a need in the art to provide improved techniques for optimizing access time when randomly accessing data on a tape storage medium.

SUMMARY OF THE PREFERRED EMBODIMENTS

[0009] Provided is a method, system, and program for storing data in a storage medium. A layout of a storage medium including a first and second user data sections is provided, wherein the first user data section comprises a faster access storage space than the second user data section. A determination is made of a first set of data to be accessed at a faster rate than a second set of data. The first set of data is written to the

[0010] In further implementations, the storage medium comprises a magnetic tape medium and the first user data section has a shorter longitudinal length than the

[0011] Further provided is a method for storing data records on a magnetic tape medium. a first set of data records is selected to write to a first user data section of the magnetic tape medium. The magnetic tape medium further includes a second user

[0012] Still further provided is a magnetic tape medium, wherein the magnetic tape medium comprises a first user data section and a second user data section. The first user data section comprises a faster access storage space than the second user data

[0013] The described implementations provide a technique to implement a fast access data section within a storage medium, such as a magnetic tape medium, to

25 allow faster random access to data in the fast access data section over data stored in
one or more additional user data sections on the storage medium.

[illegible]

20

SECRET

FIG. 1 illustrates the layout of a magnetic tape medium in a manner known in the prior art;

FIG. 3 illustrates the layout of a magnetic tape medium in accordance with implementations of the invention;

FIG. 5 illustrates logic implemented in a tape drive to access data on a magnetic storage medium in accordance with implementations of the invention; and

DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENTS

[0016] FIG. 2 illustrates an architecture of a tape storage system in which aspects of the invention are implemented. A tape cartridge 2 includes a high capacity single reel of magnetic tape 4 and a non-volatile read/writable cartridge memory 6 which maintains information about the format and layout of data on the magnetic tape. In further implementations, the tape cartridge 2 may comprise a dual wheel tape cartridge. In certain implementations, the tape cartridge 2, includes aspects of the

[illegible]

- 5 **[0017]** The tape cartridge 2 may be inserted in a tape drive 10, that includes read/write heads (not shown) capable of transferring data to and from the magnetic tape 4 in a manner known in the art. The tape drive 10 further includes one or more tape drive controllers 12 that receive Input/Output (I/O) requests from a host system 20 and is capable of executing the received I/O requests by rewinding the tape and positioning the tape heads at a particular location on the magnetic tape 4 using tape drive mechanisms and algorithms to estimate the likely location of a file on the magnet tape 4 in a manner known in the art. The tape drive 10 may be enclosed within the host system 20 or as a standalone unit or in a tape library (not shown). The tape drive 10 may connect with the host 20 through a direct interface, e.g., SCSI, 15 Fibre Channel, etc., such as the case if the tape drive 10 is connected to the host 20 or connect over a network, such as a Local Area Network (LAN), Storage Area Network (SAN), Wide Area Network (WAN), the Internet, an Intranet, etc.
- [0018]** The host 20 includes a host application 22, which may comprise a backup program, that transfers data to the tape drive 10 to sequentially write to the magnetic 20 tape 4. The host application 22 may utilize the Small Computer System Interface (SCSI) tape commands to communicate I/O requests to the tape drive 10. Details of the SCSI commands used to communicate I/O requests between the tape drive 10 and host application 22 are described in the publications “StorageSmart by IBM: Ultrium Tape Drive: SCSI Reference”, published by International Business Machines Corporation (“IBM”) as document no. WB1110-00 (August, 2000) and “SCSI-3 Stream Commands (SSC)”, published by the American Standards Institute as Working Draft, Revision 22 (Jan. 1, 2000), which publications are both incorporated 25 herein by reference in their entirety.

[0019] In accordance with SCSI tape commands, the host application 22 would write data records sequentially to the tape drive 20. To retrieve data records, the host application 22 would read data sequentially from the magnetic tape 20. To access data randomly from the magnetic tape 4, the host application 22 could send the SCSI

5 SPACE and LOCATE commands to the tape drive 20 to request a data record at an offset from the last record read from the magnetic tape 4. The host application 22 would use the SPACE command to instruct the tape drive 10 to set a new logical position relative to the current logical position, which is determined from the last data record returned by the tape drive 10. The SPACE command specifies a count field
10 indicating the number of blocks (or filemarks) to move forward (if positive) or backward (if negative). The host application 22 would use the LOCATE command to instruct the tape drive 10 to position the magnetic tape 4 to the specified logical element at the specified position.

[0020] FIG. 3 illustrates a layout format of the logical points on the magnetic tape 4 in accordance with one implementation of the invention. The layout of FIG. 3 includes sections found in the prior art LTO tape format, such as the housekeeping sections from LP1 to LP2, and the beginning (LP0 to LP1) and end (LP6 to LP7) sections. The layout of FIG. 3 further includes two separate user data sections 50 and 52, extending from LP3 to LP4 and LP3' to LP4', respectively. Data section 50 has a length that is less than the length of data section 52. Because data section 50 has a shorter longitudinal length, less time is needed to seek to a location in data section 50 during a random access operation than the time required to seek to a location in the longer data section 52. Thus, data section 50 comprises a faster data access section where the host application 22 can place data accessed more frequently, such as control and configuration data, application data, directory information, and other more frequently accessed data. Further, because the first user data section 50 is located before the second user data section 52 on the tape, data in the first user data section 50

[0021] In the prior art LTO tape format, shown in FIG. 1, each wrap extending between LP3 and LP4 in FIG. 1 would include two wrap sections of equal length. In the prior art layout of FIG. 1, the user data section between LP3 and LP4 has four bands, where each band includes six forward wraps extending from LP3 to LP4 and another six backward wraps extending from LP4 to LP3. Because each wrap includes two wrap sections, there are a total of 96 wrap sections in the user data section between LP3 and LP4 in the prior art tape layout arrangement.

section for each wrap.

[0023] FIG. 4 illustrates further details of the data structures in the cartridge memory 6, including initialization data indicating the longitudinal position of all the logical points, including LP1, LP2, LP3, LP4, LP3', LP4', LP5, LP6, and LP7. The tape drive controller 12 would use the initialization data 70 to determine the start and end of each of the user data sections 50 and 52. The cartridge memory 6 further includes a table directory 72 that includes entries for each of the wrap sections. As discussed, in implementations where there are two separate user data sections 50 and 52, the 96 possible wrap sections would be divided between these two user data sections. Thus, each wrap in the user data sections 50 and 52 would comprise a wrap section. For each of the wrap sections listed in the table directory 72, the wrap section entry may specify:

Data Set ID: specifies the Data Set Identity of the last Data Set written in this wrap section. If this wrap section does not contain valid Data Sets, then this field shall be set to (0xFFFFFFFF).

5 Record Count: If this Wrap Section is valid, this field shall contain the number of Records that are started in the current Wrap Section. If the Data Set ID of this Wrap Section is (0xFFFFFFFF) and hence this Wrap Section is invalid, the Record Count field is not defined for interchange.

10 File Mark Count: If this Wrap Section is valid, this field shall contain the number of File Marks that are within the current Wrap Section. If the Data Set ID of this Wrap Section is (0xFFFFFFFF) and hence this Wrap Section is invalid, the File Mark Count field is not defined for interchange.

CRC: This field shall specify the CRC generated for the wrap section data in the table directory 72.

15 [0024] Additional or different fields may be maintained for wrap sections in the table directory 72 to those described above.

[0025] The wrap sections would map to the lateral bands extending through the user data sections 50 and 52 in a predefined manner, such that the wrap sections comprise the forward and backward wraps within the data sections 50 and 52. The tape drive
20 controller 12 would utilize the wrap section information in the table directory 72 to estimate the longitudinal position within the wrap section of a requested data record.

Sub 24
25 [0026] In one implementation, the prior art LTO tape layout format of FIG. 1 may be modified to format the tape layout format of the described implementations of FIG. 3. For such implementations, to define the second user data section 52, LP3' may be set to a fixed value, such as 0.50 meters beyond LP4 so long as LP3' is less than LP5. The LP5 point in the prior art LTO Ultrium format (FIG. 1) then becomes LP4', and LP3 to LP5 can be 580 meters on a Type A LTO cartridge. Thus, if LP4 is set to LP3 plus 79.5 meters, then LP3 plus 79.5 meters plus 0.5 meters equals LP3 plus 80

[0027] In implementations where the magnetic tape 4 layout comprises a modification of the LTO Ultrium tape layout shown in FIG. 1, the pages in the cartridge memory 6 and the FID 56 (FIG. 3) may include code or data that would prevent tape drives that only support the standard LTO cartridge format from accessing the cartridge, so that only tape drives 10 that include logic capable of writing to the different user data sections 50 and 52 would access the magnetic tape medium 4.

[0028] FIG. 5 illustrates logic implemented in the tape drive controller 12 to access the user data sections 50 and 52. Control begins at block 100 with the tape drive controller 12 receiving I/O commands from the host 20 to reposition the magnetic tape 4 to a new position and read or write data records from that new position. The host application 22 may use the SCSI LOCATE or SPACE command to instruct the tape drive 10 to move the magnetic tape 4 to a new position. The tape drive controller 12 then processes (at block 102) the table directory 72 to determine the wrap section number n and to estimate the offset therein of the requested data set in a manner known in the art.

[0029] If (at block 106) the determined wrap section n is in data section 50, where there are wrap sections numbered 0 to 95, then the requested data set is located in the first user data section 50. In such case, the tape drive controller 12 winds (at block 108) the magnetic tape 4 the determined longitudinal offset within wrap n of the user data section 50. As discussed, when the user data sections is divided into two sections 50 and 52, the wrap section n number corresponds directly to a forward or backward wrap number within one of the bands. The tape drive controller 12 then

[0030] In the described implementations, data is written in a serpentine pattern through the user data sections 50 and 52. FIG. 6 provides a table explaining how data is written in a serpentine pattern in wrap sections through user data sections 50 and 52 in one implementation. In the serpentine pattern of FIG. 6, the first 48 wrap sections, 0-47, alternate writing forward and backward between LP3 and LP4, ending on LP3 in the 48th wrap section (wrap section 47). The next 48 wrap sections are then written in the same serpentine pattern between LP3' and LP4', alternating between forward and backward, and ending on LP3'. With the serpentine pattern implementation of FIG. 6, the tape drive must seek from LP3 to LP3', which as described above may be 80 meters apart, at the end of the first user data section 50 to the beginning of the second user data section 52.

[0031] FIG. 7 illustrates an additional serpentine pattern implementation that minimizes the distance to seek when writing data to the 49th wrap section (wrap section 48). In the serpentine pattern of FIG. 7, the tape drive writes in a serpentine pattern, alternating in the forward and backward direction between LP3 and LP4 for the first 47 wrap sections, e.g., wrap sections 0 through 46, which is the same pattern in FIG. 6 for the first 47 wrap sections. However, the pattern of FIG. 7 differs from FIG. 6 in that upon reaching the end of the 47th wrap section (the end of wrap section 46), the tape drive moves from LP4 to LP3' and then starts writing in a serpentine pattern, alternating between the forward and backward direction between LP3' and LP4' in the second user data section 52. This alternating pattern continues from wrap sections 48 through 94. There is a last possible wrap section 95, which can be written

090359-060601

10

15

25

[illegible]

- 15 Additional Implementation Details
- [0036] The preferred embodiments may be implemented as a method, apparatus or article of manufacture using standard programming and/or engineering techniques to produce software, firmware, hardware, or any combination thereof. The term “article of manufacture” as used herein refers to code or logic implemented in hardware logic
- 20 (e.g., an integrated circuit chip, Field Programmable Gate Array (FPGA), Application Specific Integrated Circuit (ASIC), etc.) or a computer readable medium (e.g., magnetic storage medium (e.g., hard disk drives, floppy disks, tape, etc.), optical storage (CD-ROMs, optical disks, etc.), volatile and non-volatile memory devices (e.g., EEPROMs, ROMs, PROMs, RAMs, DRAMs, SRAMs, firmware, programmable logic, etc.). Code in the computer readable medium is accessed and executed by a processor. The code in which preferred embodiments are implemented may further be accessible through a transmission media or from a file server over a network. In such cases, the article of manufacture in which the code is implemented

5 and that the article of manufacture may comprise any information bearing medium known in the art.

20 **[0038]** In the described implementations, there were 96 wrap sections equally divided between two different user data sections 50 and 52, where the tape drive controller 12 first writes data in a serpentine pattern to the forward and backward wraps in the first user data section 50 and then, after filling the first user data section 50, writes data in a serpentine pattern to the forward and backward wraps in the

25 second user data section 52. In additional implementations, there may be more than two user data sections, where the initialization data 70 would provide additional logical points for additional user data sections. In such cases, the 96 wrap sections may be divided among the three or more user data sections. Still further, the tape

00000000000000000000000000000000

directory and other page information in the cartridge memory may define additional or fewer wrap sections than the 96 described herein as dispersed among the multiple user data sections.

[0039] In the described implementations, data was organized into wrap sections in the user data sections. In alternative implementations, the data in the user data sections may be organized into physically consecutive groups of tracks other than wrap sections.

[0040] In the described implementations, the host application 22 communicated with the tape drive 10 using SCSI commands. In alternative implementations, different data transfer protocols other than SCSI may be used.

[0041] FIGs. 6 and 7 illustrate serpentine patterns that may be followed when writing data to wrap sections in the first 50 and second 52 user data sections. However, those skilled in the art will appreciate that different serpentine patterns may be used when writing data to the wrap sections in the user data sections 50 and 52.

15 **[0042]** In the described implementations, the directory information is stored in cartridge memory. In alternative tape implementations, such as tape cartridges that do not include a cartridge memory device, the directory information may be stored on tape or on a host system.

[0043] The described implementations provide a technique for transferring data to a
20 tape drive. Additionally, the above described logic may be used with other
input/output (I/O) devices or other storage devices, e.g., optical tape.

[0044] The described implementations utilized serpentine Linear Tape Open (LTO) technology. In alternative implementations, alternative tape technologies may be used, such as helical-scan tape drive that read/write vertical or diagonal tracks on the tape using a rotating read/write head and parallel tape drives that read/write tracks in parallel during one scan through the tape. The described implementations may be utilized with Digital Linear Tape (DLT), Quarter Inch Cassette (QIC), Travan, and any other tape technology known in the art.

[0046] The foregoing description of the preferred embodiments of the invention has been presented for the purposes of illustration and description. It is not intended to be exhaustive or to limit the invention to the precise form disclosed. Many
10 modifications and variations are possible in light of the above teaching. It is intended that the scope of the invention be limited not by this detailed description, but rather by the claims appended hereto. The above specification, examples and data provide a complete description of the manufacture and use of the composition of the invention. Since many embodiments of the invention can be made without departing from the
15 spirit and scope of the invention, the invention resides in the claims hereinafter appended.

0907E90-060801